

**Claim Amendments**

Please cancel all claims of record and substitute the new claims 115 to 157 in Enclosure

A.

**Enclosure A: Complete Listing of All the Claims in the Application**

1-114 (canceled)

115 (new): A method of scheduling on one or more processors, executions of a plurality of periodic and asynchronous processes, comprising:

(A)

automatically generating a pre-run-time schedule comprising mapping from a set of periodic process executions to a sequence of time slots on one or more processor time axes, each of the time slots having a beginning time and an end time, reserving each one of the time slots for execution of one of the periodic processes, the positions of the end time and the beginning time of each of the time slots being such that execution of the periodic processes,

including satisfaction of predetermined constraints comprising

- (1) worst-case computation times for periodic processes and asynchronous processes,
- (2) period for periodic processes,
- (3) minimum time between two consecutive requests for asynchronous processes,
- (4) deadline for periodic processes and asynchronous processes,
- (5) permitted range of offset constraints for periodic processes wherein a permitted range of offset of a periodic process comprising an interval that begins at a lower bound value and ends at an upper bound value which may be equal to the lower bound value, the duration of the time interval between the beginning of the first period of said periodic process and time zero must be greater than or equal to said lower bound value and less than or equal to said upper bound value,
- (6) precedence relations for periodic processes wherein each precedence relation being defined between a selected pair of processes comprising a first process and a second process, both said first process and said second

process being periodic processes, said first process precedes said second process, execution of said second process only allowed to start after said first process has completed its execution,

- (7) exclusion relations for periodic and asynchronous processes wherein each exclusion relation being defined between a selected pair of processes comprising a first process and a second process, said first process being either a periodic process or an asynchronous process and said second process being either a periodic process or an asynchronous process, said first process excludes said second process, no execution of said second process being allowed to occur between the time that said first process starts its execution and the time that said first process completes its computation,

can be completed between the beginning time and end time of respective time slots,

further including converting one or more asynchronous processes into corresponding new periodic processes prior to the mapping step, and mapping new periodic processes to time slots in a manner similar to mapping of other periodic processes, such that said predetermined constraints will be satisfied

(B)

during run-time using the information in the pre-run-time schedule, including the positions of the beginning time and end time of the time slots of the periodic processes, to schedule the process executions, such that said predetermined constraints will be satisfied.

116 (new): A method of scheduling on one or more processors, executions of a plurality of periodic and asynchronous processes, comprising:

(A)

automatically generating a pre-run-time schedule comprising mapping from a set of periodic process executions to a sequence of time slots on one or more processor time axes, each of the time slots having a beginning time and an end time, reserving each one of the time slots for execution of one of the periodic processes, the positions of the end

time and the beginning time of each of the time slots being such that execution of the periodic processes,

including satisfaction of predetermined constraints comprising

- (1) worst-case computation times for periodic processes and asynchronous processes,
- (2) period for periodic processes,
- (3) minimum time between two consecutive requests for asynchronous processes,
- (4) deadline for periodic processes and asynchronous processes,
- (5) permitted range of offset constraints for periodic processes wherein a permitted range of offset of a periodic process comprising an interval that begins at a lower bound value and ends at an upper bound value which may be equal to the lower bound value, the duration of the time interval between the beginning of the first period of said periodic process and time zero must be greater than or equal to said lower bound value and less than or equal to said upper bound value,
- (6) exclusion relations for periodic and asynchronous processes wherein each exclusion relation being defined between a selected pair of processes comprising a first process and a second process, said first process being either a periodic process or an asynchronous process and said second process being either a periodic process or an asynchronous process, said first process excludes said second process, no execution of said second process being allowed to occur between the time that said first process starts its execution and the time that said first process completes its computation,

can be completed between the beginning time and end time of respective time slots,

(B)

during run-time using the information in the pre-run-time schedule, including the positions of the beginning time and end time of the time slots of the periodic processes, to schedule the process executions, including allowing executions of asynchronous processes that have not been mapped to time slots in the pre-run-time schedule to be

completed within any time slot of a periodic process that has greater latitude than the latitude of the asynchronous processes in meeting their respective deadlines, such that said predetermined constraints will be satisfied.

117 (new): A method as defined in claim 115, including further executing a set of asynchronous processes that are not mapped to time slots during run-time of the processor at times which do not interfere with execution of processes mapped to time slots in the pre-run-time schedule.

118 (new): A method as defined in claim 115 including following pre-run-time scheduling and during run-time of the processor, further scheduling executions of a set of asynchronous processes that are not mapped to time slots and executions of periodic processes including said new periodic processes that are mapped to time slots such that said predetermined constraints are satisfied.

119 (new): A method as defined in claim 115, including scheduling, between the beginning time and end time of each of the time slots in the pre-run-time schedule reserved for execution of a corresponding periodic process, time capacity sufficient to complete execution of said corresponding periodic process and additional time capacity sufficient to complete execution of asynchronous processes that are not converted to new periodic processes and hence not mapped to time slots in the pre-run-time schedule and have less latitude than said corresponding periodic process in meeting their respective deadlines.

120 (new): A method as defined in claim 116, including scheduling, between the beginning time and end time of each of the time slots in the pre-run-time schedule

reserved for execution of a corresponding periodic process, time capacity sufficient to complete execution of said corresponding periodic process and additional time capacity sufficient to complete execution of asynchronous processes that are not converted to new periodic processes and hence not mapped to time slots in the pre-run-time schedule and have less latitude than said corresponding periodic process in meeting their respective deadlines.

121 (new): A method as defined in claim 115, including prior to the mapping step, automatically adjusting lengths of periods of a predetermined set of periodic processes, generating a set of reference periods, setting the length of the period of each periodic process to the length of the largest reference period that is no larger than an original period of the periodic process to form adjusted periods, and storing the adjusted periods for subsequent use in pre-run-time scheduling of executions of the periodic processes.

122 (new): A method as defined in claim 119, including prior to the mapping step, automatically adjusting lengths of periods of a predetermined set of periodic processes by generating a list of reference periods, setting the length of the period of each periodic process to the length of the largest reference period that is no larger than an original period of the periodic process to form adjusted periods, and storing the adjusted periods for subsequent use in pre-run-time scheduling of executions of the periodic processes.

123. (new) A method of determining whether an asynchronous process should or should not be converted into a new periodic process, comprising:

- (1) calculating a first processor capacity that is required for the asynchronous process if left unconverted,
- (2) calculating a second processor capacity which is required to be reserved for the new periodic process,
- (3) determining whether the ratio of said first processor capacity to said second processor capacity exceeds a predetermined threshold value.

124 (new): A method as defined in claim 115 including, prior to generating the pre-run-time schedule, determining whether each asynchronous process should or should not be converted into a new periodic process, converting a subset of a predetermined set of asynchronous processes having worst-case computation time, deadline, minimum time between two requests constraints, which have been determined to be convertible, into a set of new periodic processes having worst-case computation time, period, deadline, permitted range of offset constraints and reducing possible timing conflicts with other periodic or asynchronous processes with less latitude in meeting their deadlines, by taking into consideration the computation time requirements of the latter processes when determining the deadline of the new periodic process.

125 (new): A method as defined in claim 115 including, prior to generating the pre-run-time schedule, determining whether each asynchronous process should or should not be converted into a new periodic process, converting a subset of a predetermined set of asynchronous processes having a worst-case computation time, deadline, minimum time between two requests constraints which have been determined to be convertible, into a set of new periodic processes having release time, worst-case computation time, period, deadline, permitted range of offset constraints, wherein a permitted range of offset of each new periodic process being a subinterval of an interval or a full interval that begins at the earliest time that the corresponding being converted asynchronous process can make a request for execution, and ends at a time equal to the sum of the earliest time that said being converted asynchronous process can make a request for execution plus the period length of the new periodic process minus one time unit.

126 (new): A method as defined in claim 124, in which the determining step comprises (1) calculating a first processor capacity that is required for the asynchronous process if left unconverted,

- (2) calculating a second processor capacity which is required to be reserved for the new periodic process,
- (3) determining whether the ratio of said first processor capacity to said second processor capacity exceeds a predetermined threshold value.

127 (new): A method as defined in claim 115, including generating the pre-run-time schedule as a feasible two-part pre-run-time-schedule for execution of periodic processes that may have non-zero offsets (a) an initial part which may be of zero length, and (b) a repeating part having length which is equal to a least common multiple of lengths of the periods of the periodic processes, all executions of all periodic processes within a time interval of length equal to the length of the least common multiple of the periodic process periods being included in the repeating part of the pre-run-time schedule, wherein all said predetermined constraints being satisfied for all executions of all periodic processes within said initial part and said repeating part.

128 (new): A method as defined in claim 127, including using any offset value in a permitted range of offsets of each periodic process, including any offset value in the permitted range of offsets of any new periodic process that may have been converted from an asynchronous process, to generate said feasible pre-run-time schedule.

129 (new): A method as defined in claim 128, wherein said permitted range of offsets of any new periodic process that may have been converted from an asynchronous process is a subinterval of an interval or a full interval that begins at the earliest time that the corresponding being converted asynchronous process can make a request for execution, and ends at a time equal to the sum of the earliest time that said being converted asynchronous process can make a request for execution plus the period length of the new periodic process minus one time unit.



# BEST AVAILABLE COPY

conducting a first schedule for executions of the periodic processes within an interval starting from zero and having length equal to maximum offset value plus a bounded number of times of the length of a least common multiple of the periodic process periods, conditions for determining feasibility requiring the existence of a point in said first schedule wherein starting from the latter point the schedule repeats in subschedule interval lengths equal to a least common multiple of lengths of the periodic process periods, timing of all executions of all periodic processes within a time interval having length equal to the length of the least common multiple of the periodic process periods being included in each said repeating subschedule interval, and including satisfaction of all predetermined constraints for all executions of all periodic processes within the subschedule interval starting from time zero and ending at said point plus the length of the least common multiple of the periodic process periods in said first schedule, and checking for the first occurrence of said point in said first schedule,

(B)

generating said feasible two-part pre-run-time-schedule by

- (1) using a subschedule interval starting from time zero and ending at said point in said first schedule as said initial part of said feasible two-part pre-run-time schedule, and
- (2) using a subschedule interval starting from said point and ending at said point plus the length of the least common multiple of the periodic process periods in said first schedule as said repeating part of said feasible two-part pre-run-time schedule.

131 (new): A method as defined in claim 130, including using any offset value in a permitted range of offsets of each periodic process, including any offset value in the permitted range of offsets of any new periodic process that may have been converted from an asynchronous process, to generate said feasible pre-run-time schedule.

132 (new): A method as defined in claim 131, wherein said permitted range of offsets of any new periodic process that may have been converted from an asynchronous process is a subinterval of an interval or a full interval that begins at the earliest time that the corresponding being converted asynchronous process can make a request for execution, and ends at a time equal to the sum of the earliest time that said being converted asynchronous process can make a request for execution plus the period length of the new periodic process minus one time unit.

133 (new): A method as defined in claim 116, further comprising:

- (a) generating feasible said pre-run-time schedule for the execution of a set of hard deadline periodic processes, and of a set of soft deadline periodic processes,
- (b) assigning a criticality level and a deadline upper-limit to each soft deadline periodic process,
- (c) in the case of not finding a feasible pre-run-time schedule under said constraints , identifying a soft critical set that contains soft deadline periodic processes for which modifying the deadlines of one or more processes in said soft critical set is necessary to meet the deadlines of all hard deadline processes,
- (d) repeatedly selecting one process with a lowest criticality level among processes for which its criticality level has not reached the deadline upper-limit in the soft critical set in each case in which a feasible pre-run-time schedule has not been found, increasing the deadline of the selected process by an amount that does not exceed the deadline upper-limit of the selected process and repeatedly attempting to find a feasible pre-run-time schedule until either a feasible pre-run-time schedule is found or all the deadline upper-limits of processes in the soft critical set have been reached without finding a feasible schedule, and indicating the event of either the inability of finding a feasible schedule or of finding a feasible pre-run-time schedule, and indicating the critical set when unable to find a feasible schedule,

(e) recomputing worst-case response times of all hard deadline asynchronous processes that have not been converted to periodic processes and hence not mapped to time slots after a feasible pre-run-time schedule has been found for all hard and soft deadline periodic processes, and in every event that the worst-case response time exceeds the deadline of a particular hard deadline asynchronous process which has not been mapped to a time slot, repeatedly selecting one process that has a lowest criticality level among all soft deadline periodic processes that contribute to the worst-case response time of said particular hard deadline asynchronous process and increasing the deadline of the selected soft deadline periodic process until either (i) the worst-case response time of every hard deadline asynchronous process which has not been mapped to a time slot is less than or equal to its deadline or (ii) all the deadline upper limits of the particular set of soft deadline periodic processes which contribute to the worst-case response time of any hard deadline asynchronous process that exceeds said any hard deadline asynchronous process' deadline have been reached, and indicating the event of the latter events (i) or (ii) and the particular set.

134 (new): A method as defined in claim 115, further comprising:

- (a) during run-time, or during a pre-run-time phase, using the information in the pre-run-time schedule, including the positions of the beginning time and end time of the time slots of the periodic processes, to determine, for any point in time, whether there exists a possibility that immediate execution of a particular asynchronous process which has not been mapped to a time slot may cause the execution of any periodic process with less latitude in meeting a deadline of the latter periodic process as compared with latitude of meeting the deadline of said particular asynchronous process, to be delayed beyond a predetermined time limit, even if the periodic process is not ready for execution at said any point in time, and
- (b) during run-time, delaying execution of said particular asynchronous process if said possibility is found to exist, even if said possibility is the only reason for delaying the execution of said particular asynchronous process at said any point in time, and even if

the delay will cause the processor to be in an idle state for a time interval of non-zero length beginning from said any point in time.

135 (new): A method as defined in claim 115, further comprising:

- (a) either during run-time, or during a pre-run-time phase, using the information in the pre-run-time schedule, including the positions of the beginning time and end time of the time slots of the periodic processes, to determine, for any point in time, whether there exists a possibility that immediate execution of a particular asynchronous process which has not been mapped to a time slot may cause the execution of any periodic process with less latitude in meeting a deadline of the latter periodic process as compared with latitude of meeting the deadline of said particular asynchronous process, to be delayed beyond the end of the time slot of the periodic process in the pre-run-time schedule, even if the periodic process is not ready for execution at said any point in time, and
- (b) during run-time, delaying execution of said particular asynchronous process if said possibility is found to exist, even if said possibility is the only reason for delaying the execution of said particular asynchronous process at said any point in time, and even if the delay will cause the processor to be in an idle state for a time interval of non-zero length beginning from said any point in time.

136 (new): A method as defined in claim 115, further comprising:

- (a) either during run-time, or during a pre-run-time phase, using the information in the pre-run-time schedule, including the positions of the beginning time and end time of the time slots of the periodic processes, to determine, for any point in time, whether there exists a possibility that immediate execution of a particular asynchronous process which has not been mapped to a time slot may cause the execution of said particular asynchronous process, or the execution of some other asynchronous process, to extend beyond the beginning of the time slot of any periodic process that has not yet started in the pre-run-time schedule, even if the periodic process is not ready for execution at said any point in time, and

(b) during run-time, delaying execution of said particular asynchronous process if said possibility is found to exist, even if said possibility is the only reason for delaying the execution of said particular asynchronous process at said any point in time, and even if, the delay will cause the processor to be in an idle state for a time interval of non-zero length beginning from said any point in time.

137 (new): A method as defined in claim 115, further comprising:

(a) either during run-time, or during a pre-run-time phase, using the information in the pre-run-time schedule, including the positions of the beginning time and end time of the time slots of the periodic processes, to determine, for any point in time, whether there exists a possibility that immediate execution of a particular asynchronous process which has not been mapped to a time slot may cause the execution of any periodic process to be delayed beyond the end of the time slot of that periodic process in the pre-run-time schedule, even if the periodic process is not ready for execution at said any point in time, and

(b) during run-time, delaying execution of said particular asynchronous process if said possibility is found to exist, even if said possibility is the only reason for delaying the execution of said particular asynchronous process at said any point in time, and even if the delay will cause the processor to be in an idle state for a time interval of non-zero length beginning from said any point in time.

138 (new): A method as defined in claim 115, further comprising:

(a) either during run-time, or during a pre-run-time phase, using the information in the pre-run-time schedule, including the positions of the beginning time and end time of the time slots of the periodic processes, to determine, for any point in time, whether there exists a possibility that immediate execution of a particular first asynchronous process which has not been mapped to a time slot may cause execution of any second asynchronous process which has not been mapped to a time slot to be continuously blocked for a duration of the execution of the first asynchronous process and a duration

of the execution of any periodic process, when the second asynchronous process has less latitude in meeting the deadline of the second asynchronous process as compared with the latitude of both the first asynchronous process in meeting the deadline of the first asynchronous process and the latitude of the periodic process in meeting the deadline of the periodic process, even if neither the periodic process nor the second asynchronous process are ready for execution at said any point in time, and

(b) during run-time, delaying execution of the first asynchronous process if said possibility is found to exist, even if said possibility is the only reason for delaying the execution of the first asynchronous process at said point in time, and even if the delay will cause the processor to be in an idle state for a time interval of non-zero length beginning from said any point in time.

139 (new): A method as defined in claim 116, comprising during run-time, detecting, in a case in which no asynchronous process or periodic process that has started is to be immediately put into execution, conditions of whether there exists an execution of a first periodic process that is ready for execution and has not completed execution, and there does not exist any other execution of some second periodic process that has not yet completed, such that execution of the second periodic process is ordered before execution of the first periodic process in the pre-run-time schedule, and the time slot of the first periodic process is not nested within the time slot of the second periodic process in the pre-run-time schedule, and there does not exist any other execution of some third periodic process that is ready and has not completed execution, such that execution of the third periodic process is nested within the time slot of the first periodic process in the pre-run-time schedule, and beginning execution of the first periodic process immediately in the event said conditions are true.

140 (new): A method as defined in claim 115, including determining a worst-case response time of a particular asynchronous process which has not been mapped to a time slot that has not been converted into a periodic process using a formula comprising:

the sum of worst-case computation times of asynchronous processes and periodic processes that have less or equal latitude as compared with the latitude of said particular asynchronous process in meeting their respective deadlines, plus the maximum time that said particular asynchronous process may possibly be blocked by some asynchronous or periodic process that has greater latitude as compared with the latitude of said particular asynchronous process in meeting their respective deadlines, plus the worst-case computation time of said particular asynchronous process multiplied by the number of periodic processes with which said particular asynchronous process has an exclusion relation.

141 (new): A method as defined in claim 115, including, during a pre-run-time phase, determining by simulation a worst-case response time of a particular asynchronous process which has not been mapped to a time slot corresponding to a feasible pre-run-time schedule of periodic processes consisting of an initial part of the pre-run-time schedule which may be of zero length and a repeating part of the pre-run-time schedule, wherein for each point in time from zero to the end time of the repeating part of the run-time schedule minus one time unit, simulating execution of said particular asynchronous process using functions used to determine a run-time schedule and recording response time of said particular asynchronous process under the assumption that said particular asynchronous process arrives at a point in time under consideration, all other asynchronous processes which have not been mapped to time slots that can possibly block said particular asynchronous process arriving at one time unit prior to the point in time under consideration, and all asynchronous processes which have not been mapped to time slots that have less latitude in meeting their deadlines than latitude of said particular asynchronous process in meeting said particular asynchronous process' deadline arriving at the same said point of time under consideration, scheduling executions of periodic processes to start at the beginning time and to complete execution at the end time of their respective time slots in the pre-run-time schedule, wherein whenever said particular asynchronous process is delayed because it may block execution of some periodic process having less latitude than the latitude of said particular asynchronous process in

meeting their respective deadlines, or may block execution of some second asynchronous process for the duration of more than one execution of processes having greater latitude as compared with the latitude of the second asynchronous process, all asynchronous processes having less latitude as compared with the latitude of said particular asynchronous process in meeting their respective deadlines is delayed in order to delay said particular asynchronous process for a maximum possible amount of time, thus simulating all possible worst-case scenarios of executions of said particular asynchronous process.

142 (new): A method as defined in claim 115, including restricting every periodic process in the pre-run-time schedule to be executed strictly within its time slot.

143 (new): A method as defined in claim 115, including, during a pre-run-time phase, generating tables of safe start time intervals for the executions of asynchronous processes which have not been mapped to time slots, wherein every periodic process in the pre-run-time schedule is scheduled to be executed strictly within its time slot, wherein for every point in time of the pre-run-time schedule, it is determined whether each particular asynchronous process which has not been mapped to a time slot should be delayed, under the assumption that the actual start time of execution of every periodic process is equal to the beginning time of its time slot, and the actual end time of execution of every periodic process is equal to the end time of its time slot, wherein for every point in time of the pre-run-time schedule, in the event said particular asynchronous process is to be delayed according to the assumptions, the point in time is set to be unsafe and recorded in a corresponding entry in the table for the point in time and said particular asynchronous process.

144 (new): A method as defined in claim 115, including, during a pre-run-time phase, generating tables of safe start time intervals for the executions of asynchronous



processes which have not been mapped to time slots, wherein every periodic process is scheduled to be executed strictly within its time slot in the pre-run-time schedule, wherein for selected points in time of the pre-run-time schedule, it is determined whether each particular asynchronous process which has not been mapped to a time slot should be delayed, under the assumption that the actual start time of execution of every periodic process is equal to the beginning time of its time slot, and the actual end time of execution of every periodic process is equal to the end time of its time slot, wherein for selected points in time of the pre-run-time schedule, in the event said particular asynchronous process is to be delayed according to the assumptions, that point in time is set to be unsafe and recorded in a corresponding entry in the table for the point in time and said particular asynchronous process.

145 (new): A method as defined in claim 115 comprising during run-time, detecting at least one event of, at any point in time, whether some asynchronous process which has not been mapped to a time slot has arrived by said point in time, whether some asynchronous process or periodic process has completed its computation at said point in time, and whether said point in time is both the release time and beginning time of a time slot in the pre-run-time schedule for some periodic process, and activating a run-time scheduler at said point in time, and should said at least one event have occurred, determining whether any asynchronous process which has not been mapped to a time slot and which has arrived but has not yet been completed should be delayed or immediately put into execution, and including delaying execution of a first asynchronous process which has not been mapped to a time slot when execution of some other asynchronous process which has not been mapped to a time slot that excludes a periodic process with less or equal latitude as compared with the latitude of said first asynchronous process in meeting their respective deadlines has already started but has not yet been completed.

146 (new): A method as defined in claim 116 comprising during run-time, detecting at least one event of, at any point in time, whether some asynchronous process which has

not been mapped to a time slot has arrived by said point in time, whether some asynchronous process or periodic process has completed its computation at said point in time, and whether said point in time is both the release time and beginning time of a time slot in the pre-run-time schedule for some periodic process, and activating a run-time scheduler at said point in time, and should said at least one event have occurred, determining whether any asynchronous process which has not been mapped to a time slot and which has arrived but has not yet been completed should be delayed or immediately put into execution, and including delaying execution of a particular asynchronous process which has not been mapped to a time slot in the event there exists the possibility that immediate execution of said particular asynchronous process may cause execution of a first periodic process with less or equal latitude as compared with the latitude of said asynchronous process in meeting their respective deadlines to be delayed, when execution of the first periodic process may be preempted by execution of some second periodic process, and said particular asynchronous process cannot be preempted by the second periodic process.

147 (new): A method as defined in claim 116 comprising during run-time, detecting at least one event of, at any point in time, whether some asynchronous process which has not been mapped to a time slot has arrived by said point in time, whether some asynchronous process or periodic process has completed its computation at said point in time, and whether said point in time is both the release time and beginning time of a time slot in the pre-run-time schedule for some periodic process, and activating a run-time scheduler at said point in time, and should said at least one event have occurred, determining whether any asynchronous process which has not been mapped to a time slot that has arrived but has not yet been completed should be delayed or immediately put into execution, and including delaying execution of any asynchronous process which has not been mapped to a time slot when it is not allowed to preempt execution of any first process that has already started and has not completed execution and that excludes some other second asynchronous process which has not been mapped to a time slot which has latitude in meeting its deadline that is less than both said any first process and the latitude

of said second asynchronous process in meeting their respective deadlines, whereby blocking of the second asynchronous process by the duration of more than one execution of processes with greater latitude in meeting their respective deadlines thereof may be avoided.

148 (new): A method as defined in claim 116 comprising during run-time, detecting at least one event of, at any point in time, whether some asynchronous process which has not been mapped to a time slot has arrived by said point in time, whether some asynchronous process or periodic process has completed its computation at said point in time, and whether said point in time is both the release time and beginning time of a time slot in the pre-run-time schedule for some periodic process, and activating a run-time scheduler at said point in time, and should said at least one event have occurred, determining whether any asynchronous process which has not been mapped to a time slot and which has arrived but has not yet been completed should be delayed or immediately put into execution, and including delaying execution of a particular asynchronous process which has not been mapped to a time slot to allow preemption of execution of said particular asynchronous process by execution of a first periodic process that is ready for execution and that has a latitude which is less than or equal to the latitude of said particular asynchronous process in meeting their respective deadlines, when said particular asynchronous process does not exclude the first periodic process and does not exclude any other asynchronous process which has not been mapped to a time slot with a latitude that is less than the latitude of the first periodic process in meeting their respective deadlines, and there does not exist execution of some second periodic process that has not been completed such that execution of the second periodic process is ordered before execution of the first periodic process and execution of the time slot of the first periodic process is not nested within the time slot of the second periodic process in the pre-run-time schedule.

149 (new): A method as defined in claim 115, including, during a pre-run-time phase, determining by simulation a worst-case response time of a particular asynchronous process which has not been mapped to a time slot corresponding to a feasible pre-run-time schedule of periodic processes, wherein for each point in time from zero to the end time of the repeating part of the run-time schedule minus one time unit, simulating execution of said particular asynchronous process using functions used to determine a run-time schedule and recording response time of said particular asynchronous process under the assumption that said particular asynchronous process arrives at a point in time under consideration, all other asynchronous processes which have not been mapped to time slots that can possibly block said particular asynchronous process arriving at one time unit prior to the point in time under consideration, and all asynchronous processes which have not been mapped to time slots that have less latitude than latitude of said particular asynchronous process in meeting their respective deadlines arriving at the same said point of time under consideration, scheduling executions of periodic processes to start at the beginning time and to complete execution at the end time of their respective time slots in the pre-run-time schedule, simulating all possible worst-case scenarios of executions of said particular asynchronous process.

150 (new): A method of scheduling on one or more processors, executions of a plurality of periodic and asynchronous processes, comprising:

(A)

automatically generating a pre-run-time schedule comprising mapping from a set of periodic process executions to a sequence of time slots on one or more processor time axes, each of the time slots having a beginning time and an end time, reserving each one of the time slots for execution of one of the periodic processes, the positions of the end time and the beginning time of each of the time slots being such that execution of the periodic processes,

including satisfaction of predetermined constraints comprising

- (1) worst-case computation times for periodic processes and asynchronous processes,
- (2) period for periodic processes,

- (3) minimum time between two consecutive requests for asynchronous processes,
- (4) deadline for periodic processes and asynchronous processes,
- (5) precedence relations for periodic processes wherein each precedence relation being defined between a selected pair of processes comprising a first process and a second process, both said first process and said second process being periodic processes, said first process precedes said second process, execution of said second process only allowed to start after said first process has completed its execution,
- (6) exclusion relations for periodic and asynchronous processes wherein each exclusion relation being defined between a selected pair of processes comprising a first process and a second process, said first process being either a periodic process or an asynchronous process and said second process being either a periodic process or an asynchronous process, said first process excludes said second process, no execution of said second process being allowed to occur between the time that said first process starts its execution and the time that said first process completes its computation,

can be completed between the beginning time and end time of respective time slots, including scheduling, between the beginning time and end time of each of the time slots in the pre-run-time schedule reserved for execution of a corresponding periodic process, time capacity sufficient to complete execution of said corresponding periodic process and additional time capacity sufficient to complete execution of asynchronous processes which are not converted to new periodic processes and hence not mapped to time slots in the pre-run-time schedule and which have less latitude than said corresponding periodic process in meeting their respective deadlines.

(B)

during run-time using the information in the pre-run-time schedule, including the positions of the beginning time and end time of the time slots of the periodic processes, to schedule the process executions, including allowing executions of asynchronous processes that have not been mapped to time slots in the pre-run-time schedule to be completed within any time slot of a periodic process that has greater latitude than the

latitude of the asynchronous processes in meeting their respective deadlines, such that said predetermined constraints will be satisfied.

151 (new): A method as defined in claim 150 including, prior to the mapping step, converting one or more asynchronous processes having a worst-case computation time, deadline, minimum time between two requests constraints, into a set of new periodic processes, and reducing possible timing conflicts with other periodic or asynchronous processes with less latitude in meeting their respective deadlines, by taking into consideration the computation time requirements of the latter processes when determining the deadline of each of the new periodic processes, and mapping the new periodic process in a manner similar to mapping of other periodic processes.

152 (new): A method as defined in claim 150 including, prior to generating the pre-run-time schedule, determining whether each asynchronous process should or should not be converted into a new periodic process, converting a subset of a predetermined set of asynchronous processes having worst-case computation time, deadline, minimum time between two requests constraints, which have been determined to be convertible, into a set of new periodic processes having worst-case computation time, period, deadline, and reducing possible timing conflicts with other periodic or asynchronous processes with less latitude in meeting their respective deadlines, by taking into consideration the computation time requirements of the latter processes when determining the deadline of the new periodic process.

153 (new): A method of scheduling on one or more processors, executions of a plurality of periodic processes, comprising:

(A)

automatically generating a pre-run-time schedule comprising mapping from a set of periodic process executions to a sequence of time slots on one or more processor time

axes, each of the time slots having a beginning time and an end time, reserving each one of the time slots for execution of one of the periodic processes, the positions of the end time and the beginning time of each of the time slots being such that execution of the periodic processes,

including satisfaction of predetermined constraints comprising

- (1) worst-case computation times,
- (2) period,
- (3) deadline,
- (4) permitted range of offset constraints wherein a permitted range of offset of a periodic process comprising an interval that begins at a lower bound value and ends at an upper bound value which may be equal to the lower bound value, the duration of the time interval between the beginning of the first period of said periodic process and time zero must be greater than or equal to said lower bound value and less than or equal to said upper bound value,
- (5) precedence relations wherein each precedence relation being defined between a selected pair of processes comprising a first process and a second process, both said first process and said second process being periodic processes, said first process precedes said second process, execution of said second process only allowed to start after said first process has completed its execution,
- (6) exclusion relations wherein each exclusion relation being defined between a selected pair of processes comprising a first process and a second process, said first process being either a periodic process or an asynchronous process, said first process excludes said second process, no execution of said second process being allowed to occur between the time that said first process starts its execution and the time that said first process completes its computation,

can be completed between the beginning time and end time of respective time slots,

(B)

during run-time using the information in the pre-run-time schedule, including the positions of the beginning time and end time of the time slots of the periodic processes, to schedule the process executions such that said predetermined constraints will be satisfied.

154 (new): A method as defined in claim 153, including generating the pre-run-time schedule as a feasible two-part pre-run-time-schedule for execution of periodic processes that may have non-zero offsets (a) an initial part which may be of zero length, and (b) a repeating part having length which is equal to a least common multiple of lengths of the periods of the periodic processes, all executions of all periodic processes within a time interval of length equal to the length of the least common multiple of the periodic process periods being included in the repeating part of the pre-run-time schedule, wherein all said predetermined constraints being satisfied for all executions of all periodic processes within said initial part and said repeating part.

155 (new): A method of determining whether each asynchronous process in an original set of asynchronous processes should or should not be converted into a new periodic process and included in a set of periodic processes that will be mapped to time slots in a pre-run-time schedule,

each asynchronous process in said original set of asynchronous processes has predetermined asynchronous process constraints comprising worst-case computation time, deadline, and minimum time between two consecutive requests constraints,

each periodic process in said set of periodic processes has predetermined periodic process constraints comprising permitted range of offset, worst-case computation time, deadline, period constraints,

each new periodic process has new periodic process constraints comprising permitted range of offset, worst-case computation time, deadline, period constraints, said method steps comprising:

(A)



selecting one asynchronous process in said original set of asynchronous processes,

- (1) tentatively converting said one asynchronous process into a corresponding new periodic process, such that said asynchronous process constraints of said one asynchronous process will be satisfied by said new periodic process constraints of said corresponding new periodic process,
- (2) calculating a first processor capacity which is required to be reserved if said one asynchronous process is not converted,
- (3) calculating a second processor capacity which is required to be reserved if said one asynchronous process is converted,
- (4) including a copy of said one asynchronous process in a second set of asynchronous processes if the ratio of said first processor capacity to said second processor capacity exceeds a predetermined threshold, otherwise including said tentatively converted corresponding new periodic process in said set of periodic processes,

(B)

repeating A until every asynchronous processes in said original set of asynchronous processes has been selected,

(C)

repeating B,

including removing from the periodic set any corresponding new periodic process that was included in the periodic set in any previous execution of B whenever a copy of a asynchronous process is to be included in the second set of asynchronous processes in the current execution of step B, further including removing any copy of an asynchronous process from the second asynchronous set that was included in the second asynchronous set in any previous execution of step A whenever any corresponding new periodic process is to be included in the set of periodic processes in the current execution of step A,

(D)

terminating when no more changes to the second asynchronous set have occurred between the beginning and the end of any execution of C.

156 (new): A method as defined in claim 152, in which the determining step comprises

- (1) calculating a first processor capacity that is required for the asynchronous process if left unconverted,
- (2) calculating a second processor capacity which is required to be reserved for the new periodic process,
- (3) determining whether the ratio of said first processor capacity to said second processor capacity exceeds a predetermined threshold value.

157 (new): A method as defined in claim 154, further including

(A)

constructing a first schedule for executions of the periodic processes within an interval starting from zero and having length equal to maximum offset value plus a bounded number of times of the length of a least common multiple of the periodic process periods, conditions for determining feasibility requiring the existence of a point in said first schedule wherein starting from the latter point the schedule repeats in subschedule interval lengths equal to a least common multiple of lengths of the periodic process periods, timing of all executions of all periodic processes within a time interval having length equal to the length of the least common multiple of the periodic process periods being included in each said repeating subschedule interval, and including satisfaction of all predetermined constraints for all executions of all periodic processes within the subschedule interval starting from time zero and ending at said point plus the length of the least common multiple of the periodic process periods in said first schedule, and checking for the first occurrence of said point in said first schedule,

(B)

generating said feasible two-part pre-run-time-schedule by

- (1) using a subschedule interval starting from time zero and ending at said point in said first schedule as said initial part of said feasible two-part pre-run-time schedule, and

(2) using a subschedule interval starting from said point and ending at said point plus the length of the least common multiple of the periodic process periods in said first schedule as said repeating part of said feasible two-part pre-run-time schedule.